Segmentation of Coarse and Fine Scale Features Using Multi-scale Diffusion and Mumford-Shah

Jeremy D. Jackson[1], Anthony Yezzi, Jr.[1,⋆],
Wes Wallace[2], and Mark F. Bear[2,⋆⋆]

[1] School of Electrical and Computer Engineering, Georgia Institute of Technology
Atlanta, GA 30332
gtg120d@prism.gatech.edu
[2] Department of Neuroscience, Howard Hughes Medical Institute/Brown University
Providence, RI 02912
wwallace@brown.edu

Abstract. Here we present a segmentation algorithm that uses multi-scale diffusion with the Mumford-Shah model. The image data inside and outside a surface is smoothed by minimizing an energy functional using a partial differential equation that results in a trade-off between smoothing and data fidelity. We propose a scale-space approach that uses a good deal of diffusion as its coarse scale space and that gradually reduces the diffusion to get a fine scale space. So our algorithm continually moves to a particular diffusion level rather than just using a set diffusion coefficient with the Mumford-Shah model. Each time the smoothing is decreased, the data fidelity term increases and the surface is moved to a steady state. This method is useful in segmenting biomedical images acquired using high-resolution confocal fluorescence microscopy. Here we tested the method on images of individual dendrites of neurons in rat visual cortex. These dendrites are studded with dendritic spines, which have very small heads and faint necks. The coarse scale segments out the dendrite and the brighter spine heads, while avoiding noise. Backing off the diffusion to a medium scale fills in more of the structure, which gets some of the brighter spine necks. The finest scale fills in the small and detailed features of the spines that are missed in the initial segmentation. Because of the thin, faint structure of the spine necks, we incorporate into our level set framework a topology preservation method for the surface which aids in segmentation and keeps a simple topology.

1 Introduction

Global segmentation algorithms have the benefit of being able to extract an object and its prominent features from an image or image volume. They have this capability because they segment an image based on properties such as average pixel intensity of a region or differing textures of regions. Some of these methods

⋆ Supported by NSF grant CCR-0133736 and NIH grant R01-HLS0004-01A1.
⋆⋆ This work is supported in part by a grant from the National Eye Institute.

L.D. Griffin and M. Lillholm (Eds.): Scale-Space 2003, LNCS 2695, pp. 615–624, 2003.

are detailed in [2], [5],[7], [8],[9]. While a global perspective avoids the noise that more edge-based detectors would get caught up on, it can lose the fine scale features of the object in capturing a coarse estimate of the object. It would be useful to have a method with which to include some of these finer scale features after this coarse segmentation has been done. In our work, a method of this type was necessary to solve the problem of segmenting a topographically complex biological structure from a three dimensional image volume.

The structures in question are the dendrites of pyramidal neurons in rat visual cortex. These dendrites are studded with individual tiny branchlets called spines. The spines are sites of synaptic contact between neurons, and their 3-d morphology is thought to be a marker of the functional state of individual synapses. The fine structure of spines has been extensively investigated at the electron-microscope level – they are known to be bulbous in shape and always connected to the dendrite by very thin necks (with diameter on the order of 0.1 micron). [10]

We obtained 3D image volumes of spiny dendrites as follows: pyramidal neurons in fixed tissue slices of rat visual cortex were intracellularly injected with the fluorescent dye Alexa-488 (Molecular Probes Inc., Eugene, OR; emission peak = 517 nm). Individual dendritic segments were imaged in 3D using an olympus fluoview confocal microscope, at zoom factor 8, with a 63x NA 1.2 water-immersion lens. The voxel size of these images was 0.09 x 0.09 x 0.15 microns (actually slightly above the diffraction limit of this imaging system). 3D images were preprocessed using simple operations to improve contrast and reduce noise. Images were then deconvolved using an adaptive blind deconvolution algorithm (Autoquant Imaging, Watervliet, NY).

In these images, the dendrite is more brightly fluorescent than the spines, due to the greater volume of fluorescent dye it contains. The spine necks in particular can be very faint both because of their very small volume, and because their size is at the limit of resolution of the confocal microscope. Some of the spine heads are dim as well. This is apparent by looking at a full 2D slice of the 3D images in Fig. 1 and a close up of a section of the dendrite and its spines in Fig. 2.

The regional methods only capture the dendrite and some of the spine heads. The first step to solve this problem is to set a smoothing parameter in the Mumford-Shah segmentation method so that it becomes a regional algorithm that gives a coarse segmentation of the dendrite. Then this smoothing term is gradually reduced to capture some fine scale features. It is this stepping down of the diffusion term that gradually gets a correct segmentation of the spine heads and the necks that connect them to the dendrite.

2 The Mumford Shah Model

Here in this section we present the variational formulation of the main segmentation algorithm (a multi-scale version of Mumford-Shah) that was used in this project. This algorithm was implemented in a level set formulation according to [11]. Other level set implementations of Mumford-Shah are in [1], [6] and the

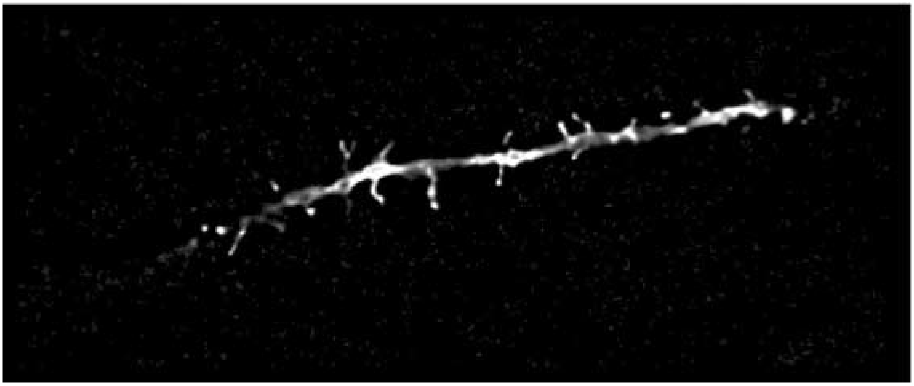

Fig. 1. 2D image plane from the middle of a 3D volume, showing the dendrite with spines branching off

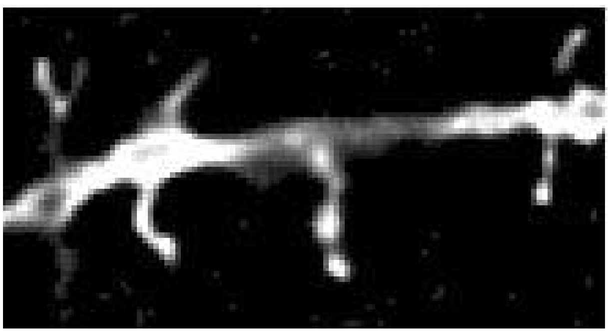

Fig. 2. Closeup of Fig. 1, showing the fine structure of spines. Pixel size = 0.09 x 0.09 microns

model itself is in [4]. The level set is evolved using a PDE that minimizes a given energy functional. More about level set theory can be found in [12].

To implement the Mumford and Shah model, a smooth estimate of the foreground and one of the background is needed so there can exist a piecewise smooth estimate of the image data with the surface being the discontinuity between the two estimates. Based on these smooth estimates, the level set which contains the surface (in this case, a three dimensional surface) is evolved to minimize the following energy functional:

$$E = \alpha \iiint_R (I-f)^2 dV + \alpha \iiint_{R^c} (I-g)^2 dV + \beta \iiint_R |\nabla f|^2 dV + \beta \iiint_{R^c} |\nabla g|^2 dV + \gamma \iint_S d\sigma \quad (1)$$

where I is the image volume, f is the smooth estimate of the image in the foreground R, g is the smooth estimate of the image in the background R^c, and S is the surface. The first two terms in the energy functional are data fidelity terms that make sure that the smooth estimates of the foreground and background match the image data as much as possible. The next two terms keep the norm squared gradients of the smooth estimates f and g as small as possible which results in a smoother f and g. The last term of the energy functional is used to penalize surface area. The parameters $\alpha, \beta, \gamma \in [0,1] \subset \Re$ should all add up to 1 so they can be used as weights to either increase data fidelity or smoothness or penalizing of surface area. So then the level set is evolved according to the flow

$$\phi_t = -\alpha((I-g)^2 - (I-f)^2)\boldsymbol{N} + \beta(|\nabla g|^2 - |\nabla f|^2)\boldsymbol{N} + \gamma\kappa\boldsymbol{N} \tag{2}$$

where $\boldsymbol{N}$ is the inward normal of the surface S. The derivation of this can be found in [1] and [4]. With each evolution of the level set ϕ, we need to get the new smooth estimates f and g. This is done using the same energy functional as above but minimization is done with respect to f when evolving the smooth estimate f. Using the Calculus of Variations, the first variation is used to get the Euler-Lagrange equations necessary to evolve the smooth function to a steady state based on the the energy functional. The resulting equation to evolve the smooth function f is

$$f_t = 2(\alpha(I-f) + \beta \Delta f) \tag{3}$$

where Δf is the laplacian of f:

$$\Delta f = f_{xx} + f_{yy} + f_{zz}. \tag{4}$$

Evolving g is similar.

A piecewise constant version of this is given by Chan and Vese in [5]. The energy functional is given by:

$$E = \alpha \iiint_R (I-u)^2 dV + \alpha \iiint_{R^c} (I-v)^2 dV + \gamma \iint_S d\sigma. \tag{5}$$

where u and v are the means inside and outside the surface respectively. The evolution of the the level set is given by

$$\phi_t = -\alpha(u-v)(I-u+I-v)\boldsymbol{N} + \gamma\kappa\boldsymbol{N}. \tag{6}$$

The Chan-Vese flow can also be looked at as the (β=∞) case (total smoothing) of Mumford-Shah.

3 Multi-scale Diffusion with Mumford-Shah

So in our algorithm, the coarse Mumford-Shah segmentation that we begin with is the ($\beta = \infty$) case which is equivalent to the Chan-Vese piecewise constant model. We evolve the Mumford-Shah flow to steady state, decrease the smoothing parameter and increase the data fidelity term.

First let us see why the ($\beta = \infty$) case is our coarse scale space which will segment prominent features of the image only. The update of the level set can be rearranged as such:

$$\phi_t = -2\alpha(u - v)(I - \frac{u+v}{2})\boldsymbol{N} + \gamma\kappa\boldsymbol{N}. \quad (7)$$

If the surface is initialized so that it is outside of the object we want to segment, then the term $-2\alpha(u - v)$ should not change sign while the surface is evolving. The term $I - \frac{u+v}{2}$ shows us that the flow will move the surface according to u (the mean of the image data inside the surface) and v (the mean of the image data outside the surface) so that the energy

$$E = \alpha \iiint_R (I - u)^2 dV + \alpha \iiint_{R^c} (I - v)^2 dV \quad (8)$$

is as small as possible. So what happens is with each iteration the means are computed and the surface moves past a pixel in I if it is less than $\frac{u+v}{2}$. This is the case if we ignore the surface area penalty which gets rid of bright pieces of noise because they have high curvature. The value $\frac{u+v}{2}$ in this case can be looked at as a threshold that gets larger as the surface segments a bright object. This flow gives a segmentation of all of the very prominent features of the object. The problem with this is the single value $\frac{u+v}{2}$ that is used to move the surface at all points in the image. This tends to skip over fine detail that might be fainter than most of the rest of the object. In the case of dendrites, the main dendrite and the head of the spines are segmented very well, but the dimmer spine necks are totally skipped over.

To fix this we need the Mumford-Shah flow (with $\beta \neq \infty$) which uses a value $\frac{f+g}{2}$ to decide whether to pass by a pixel or not. Since f and g are smooth functions, there is a more adaptive threshold that passes by pixels depending on a value that is more local to the pixel since f or g at each pixel is smoothed out by its neighboring pixels. This is preferred over a global smoothing ($\beta = \infty$) which results in $f = u$ and $g = v$. This allows Mumford-Shah to capture some of the fine detail. So the premise of our algorithm is to keep backing off the smoothing to acquire more and more detail of the object from a very nice, but rough initial estimate. This gradual aquiring of features in a multi-step fashion allows the flow to accurately capture more detail than a Mumford-Shah flow with a set diffusion. The set Mumford-Shah flow does not get these details as well as the multi-step version because it has no good coarse segmentation to build upon.

Also an assumption that we made in the segmentation of dendrites is that the background is constant (fairly close to zero) which turns out to be true for all the data we have worked on. This allows us to use v or zero as the estimate for the background which speeds up the process since it is not necessary to use a PDE to find the smooth function g each time the surface needs to be evolved.

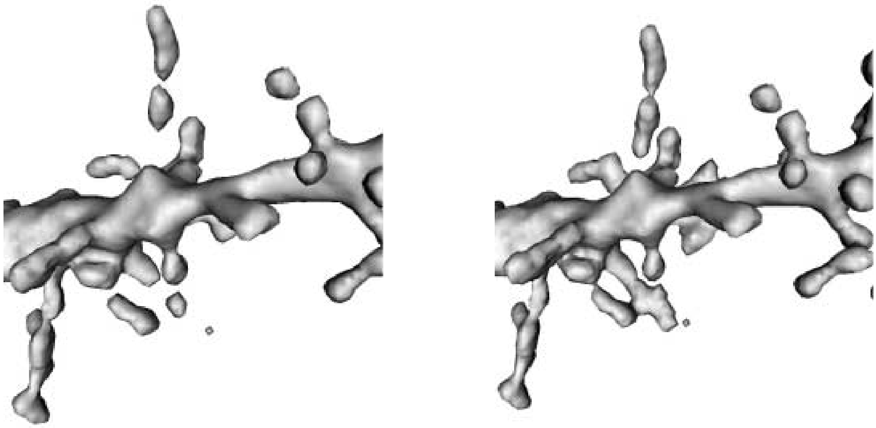

Fig. 3. Set Mumford-Shah vs.Multi-Scale Mumford-Shah

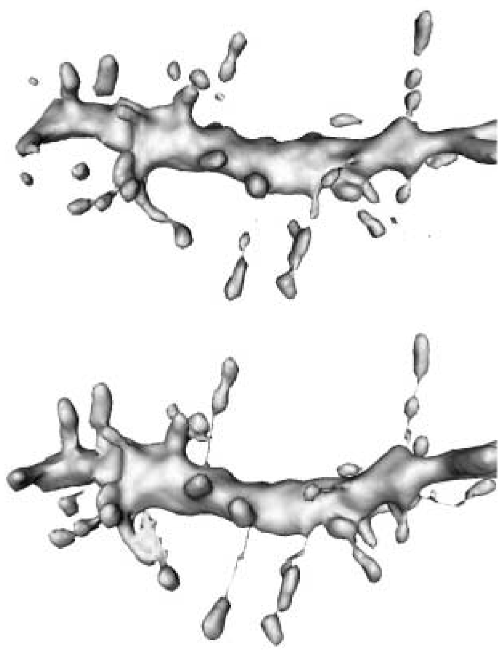

Fig. 4. Mumford-Shah ($\beta = \infty$) case: no topology preservation vs. topology preservation

4 Topology Preservation with Mumford Shah

It would be nice to keep objects that we segment to be as realistic as possible. In the case with dendrites there are no holes of any kind; so a dendrite should be

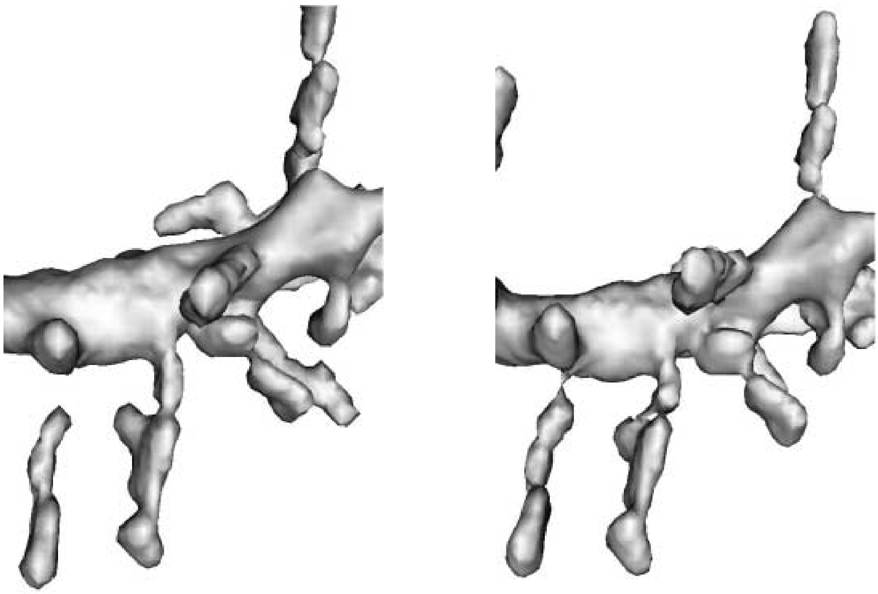

Fig. 5. Connection example: Multi-Scale Mumford-Shah without and with topology preservation

topologically equivalent to a sphere. To keep this realism in our surface we need a flow that preserves the topology of the surface. This will also preserve fine-scale features (i.e. the necks of the dendritic spines). We use the method in [14] which preserves the topology of a surface in a level-set methodology. This method looks for *simple points* as described in [15], [16], and [17]. If this preservation is not done, the surface will pinch off the necks and just segment the dendrite and the spine heads and will not have a simple topology.

Our level set function uses values below zero to denote the inside of the surface (the zero level set) and values above zero to denote pixels that are outside the surface. When a value of our level set ϕ wants to change sign, i.e. a pixel wants to change from foreground to the background or vice versa, it is possible that the change will cause a change in topology. To keep this from happening we look at a point in the level set when it is going to change sign. If this will cause a break in topology (the point is not a *simple point*), we just set the value of the level set at that point to be some small number ϵ that has the same sign as the point had before.

This topology preservation helps at each step in the evolution of our surface. The initial ($\beta = \infty$) Mumford-Shah flow needs to have this preservation so that it will not break topology so our initial coarse estimate is still topologically equivalent to a sphere. If this topology preservation is not in place, the necks of the dendritic spines would get pinched off as shown in Fig. 4.

It is possible to get these necks back without doing topology preservation and just running the multi-step Mumford-Shah. The advantage of having topology

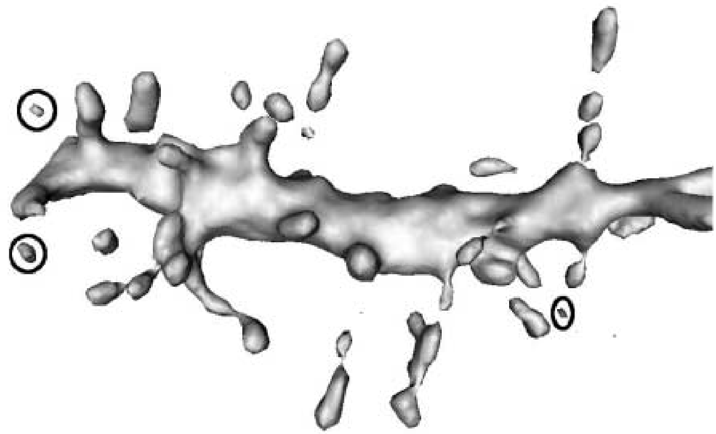

Fig. 6. Noise example: Mumford-Shah ($\beta = \infty$) case without topology preservation

Fig. 7. Noise example: Mumford-Shah ($\beta = \infty$) case with topology preservation

preservation is that there is a piece of surface that is already connecting the spine head and the dendrite where the neck should be. This makes it easier for the multi-step Mumford-Shah to expand out over that neck. Whereas without the neck surface there, the neck does get found, but in the case of a totally missing neck or extremely faint data the multi-scale Mumford-Shah will not fill in the neck completely and so it will not totally connect the spine head to the dendrite.

Another benefit of having the topology preservation is that it helps get rid of pieces of noise. With topology preservation, there is some surface that connects the noise to the dendrite. Without the surface connecting the noise to the dendrite, the noise has its own local smooth function and the areas near it have a

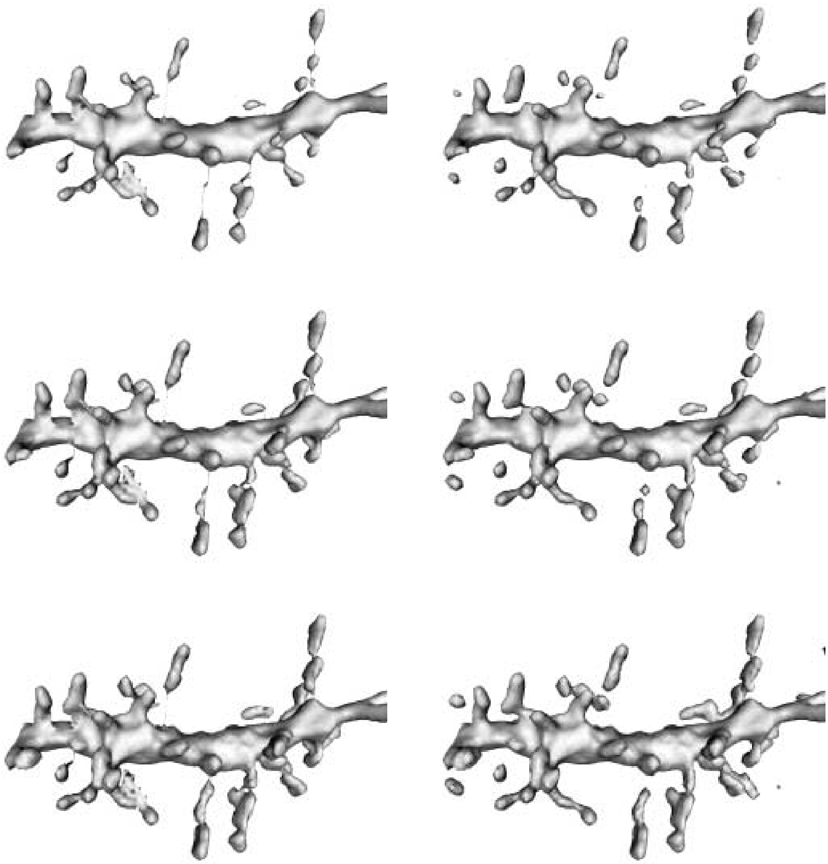

Fig. 8. Progression of Multi-Scale Mumford-Shah with and without topology preservation

smooth function close to or equal to zero because they are background. Therefore the noise has a greater chance of staying in the segmentation as shown in Figure 6. With a surface connecting the noise to the dendrite, the smooth function close to that region of noise will be higher causing the noise to disappear as in Figure 7.

5 Conclusion

Here we have shown a method to segment fine scale features of a biological object. This scale-space approach of a multi-scale Mumford-Shah is very good for capturing coarse and then fine-scale features. Also the preservation of topology allows for a more realistic segmentation with no breaks in topology. In the case of dendritic spines, we have prior knowledge of their topology and therefore we can require that spine heads remain connected to the dendrite by a neck, even when there is no data for a neck. Topology preservation also improves segmentation of spine necks in cases where the data for the neck exists but is very faint. This is evident in Figure 8 where the progression of Multi-scale Mumford-Shah is shown with and without topology preservation. The Multi-scale Mumford-Shah with topology preservation captures the dendrite quite well.

References

1. Tsai, A., Yezzi A., Wilsky A.: Curve Evolution Implementation of the Mumford-Shah Functional for Image Segmentation, Denoising, Interpolation, and Magnification. IEEE Trans. on Image Processing (2001) 1169–1184
2. Yezzi, A., Tsai, A., Wilsky, A.: A statistical approach to snakes for bimodal and trimodal imagery. Int. Conf. on Computer Vision. (1999) **2** 898–903
3. Yezzi, A.: Modified curvature motion for image smoothing and enhancement. IEEE Trans. Image Processing **7** (1998) 345–352
4. Mumford D., Shah J.: Optimal approximations by piecewise smooth functions and associated variational problems. Commun. Pure Appl. Math (1989)
5. Chan, T.F, Vese, L.A.: Active contours without edges. IEEE Trans. Image Processing **10** (2001) 266-277
6. Chan, T.F, Vese, L.A.: A level set algorithm for minimizing the Mumford-Shah functional in image processing. IEEE Proc. on Variational and Level Set Meth. in Comp. Vision (2001) 161–168
7. Paragios, N., Deriche, R.: Geodesic active regions for texture segmentation. INRIA, France, Res. Rep. 3440 (1998)
8. Ronfard, R.: Region-based strategies for active contour models. Int. J. Comput. Vis. **13** (1994) 229–251
9. Zhu, S., Yuille, A.: Region competition: Unifying snakes, region growing, and Bayes/MDL for multiband image segmentation. IEEE Trans. Pattern Anal. Machine Intell. **18** (1996) 884–900
10. Nimchinsky E. A., Sabatini B.L. and Svoboda K. (2002). "Structure and function of dendritic spines." Annu. Rev. Physiol. 64: 313-353.
11. Osher, S., Sethian, J.: Fronts propagation with curvature dependent speed: Algorithms based on Hamilton-Jacobi formulations. J. Comput. Physics **79** (1988) 12–49
12. J.A. Sethian: Level Set Methods and Fast Marching Methods: Evolving Interfaces in Geometry, Fluid Mechanics, Computer Vision, and Material Science. Cambridge, U.K.:Cambridge University Press (1999)
13. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active contour models. Int. J. Comput. Vis **1** (1987) 321-331
14. Han, X., Xu, C., Tosun, D., Prince, J.L.: Corical Surface Reconstruction Using a Topology Preserving Geometric Model. IEEE Trans. on Medical Imaging (2002) 109-121
15. Malandain, G., Bertrand G.: Fast Characterization of 3D Simple Points. IEEE Pattern Recognition (1992) 232–235
16. Malandain, G., Bertrand G.: A new characterization of three-dimensional simple points. Pattern Recognition Letters **15** (1994) 169–175
17. Bertrand, G.: Simple points, toplogical numbers and geodesic negihborhoods in cubic grids. Pattern Recognition Letters **15** (1994) 1003–1011
18. Weinstock, R.: Calculus of Variations: With Applications to Physics and Engineering. New York:Dover Pub. Inc. (1974)